

WHAT IS CLAIMED IS:

1. A multiprocessor computer system, comprising:

a plurality of nodes, each node including:

5 an interface to a local memory subsystem, the local memory subsystem storing a multiplicity of memory lines of information and a directory;

a memory cache for caching a multiplicity of memory lines of information, including memory lines of information stored in a remote memory subsystem that is local to another node;

10 the directory including an entry associated with a memory line of information stored in the local memory subsystem, the entry including an identification field for identifying a subset of nodes from the plurality of nodes caching the memory line of information;

15 the identification field configured to comprise a plurality of bits at associated positions within the identification field;

a protocol engine implementing a cache coherence protocol, said protocol engine configured to

20 associate with each respective bit of the identification field one or more nodes of the plurality of nodes, including a respective first node, wherein the one or more nodes associated with each respective bit are determined by reference to the position of the respective bit within the identification field;

25 set each bit in the identification field of the directory entry associated with the memory line for which the memory line is cached in at least one of the associated nodes;

send an initial invalidation request to no more than a first predefined number of the nodes associated with set bits in the identification field of the directory entry associated with the memory line.

2. The system of claim 1, wherein the protocol engine is further configured to:

store in the identification field of the directory entry associated with the memory line one or more node identifiers that identify a subset of the plurality of nodes in which the memory line is cached, when the subset of nodes includes fewer than a second predefined number of nodes, wherein each node identifier stored in the identification field occupies a plurality of the bits of the identification field; and

send an initial invalidation request to no more than the first predefined number of the nodes whose node identifiers are stored in the identification field of the directory entry associated with the memory line.

3. The system of claim 2, wherein the protocol engine is configured to

respond to a request to share the memory line with an additional node such that the memory line will be cached in more than the second predefined number of nodes, by reconfiguring the identification field of the directory entry associated with the memory line by setting each bit in the identification field that is associated with any of the nodes in which the memory line is cached or will be cached upon servicing the request.

4. The system of claim 2, wherein

the directory entry further includes a state field, the state field indicating whether the identification field is configured to store the one or more node identifiers, the state field otherwise indicating that the identification field is configured to associate with each respective bit of the identification field one or more nodes of the plurality of nodes.

5. The system of claim 1, wherein

the plurality of nodes is a first number of nodes;

the plurality of bits in the identification field is a second number of bits;

the first number is greater than the second number; and

the protocol engine is configured to associate both the first respective node and a respective second node with a particular one of the bits of the identification field, to generate a first node identifier corresponding to the first respective node in accordance with the position of the particular bit within the identification field, and to generate a second node identifier corresponding to the second respective node in accordance with the position of the particular bit within the identification field and the second number of bit in the identification field.

6. The system of claim 1, wherein

the protocol engine is configured to send the initial invalidation request to the first node associated with a particular set bit in the identification field.

7. The system of claim 1, wherein

the identification field is subdivided to form a number of groups of bits equal to the first predefined number; and

the protocol engine is configured to send at most one invalidation request for each group of bits, wherein the at most one invalidation request for each group of bits is sent to a first node, if any, associated with a set bit in the group of bits.

8. The system of claim 7, wherein the first predefined number is at least four and no greater ten.

9. The system of claim 7, wherein

the protocol engine is configured to include in the initial invalidation request sent to the first node associated with one of the groups of bits in the identification field a pattern of bits based on the one group of bits in the identification field, such that a recipient node of the initial invalidation request can derive from the pattern of bits a next recipient node, if any, to which to send a second invalidation request corresponding to the initial invalidation request.

10. The system of claim 7, wherein

the protocol engine is configured send a respective version of the initial invalidation request to the first node, if any, associated with each group of bits of the identification field, and to include in each respective version of the initial invalidation request a pattern of bits based on the respective group of bits in the identification field, such that each first node can derive a next recipient node, if any, from the pattern of bits in the respective version of the initial invalidation request received by the first node, wherein the next recipient node is to be sent a second invalidation request corresponding to the initial invalidation request.

11. The system of claim 1, wherein

the protocol engine is configured to include in the initial invalidation request a pattern of bits based on at least a subset of the plurality of bits in the identification field, such that a recipient node of the initial invalidation request can derive from the pattern of bits a next recipient node, if any, to which to send a second invalidation request corresponding to the initial invalidation request.

12. The system of claim 1, wherein

the identification field is subdivided to form a number of groups of bits; and the protocol engine is configured to send, for each respective group of bits, the initial invalidation request to a first node, if any, associated with a set bit in the respective group of bits, and to include in the invalidation request a pattern of bits based on the respective group of bits in the identification field.

13. The system of claim 12, wherein

the protocol engine is configured to send the initial invalidation request to a second node associated with a set bit in the respective group of bits when the first node is a node requesting exclusive ownership of the memory line of information.

14. The system of claim 12, wherein

the protocol engine is configured to send the initial invalidation request to a second node associated with a set bit in the respective group of bits when the first node is a home node of the memory line of information.

15. The system of claim 1, wherein

the protocol engine is further configured to forward an invalidation request received by the protocol engine to a next node identified in the invalidation request.

16. The system of claim 1, wherein

5 the protocol engine is further configured to associate with a respective subset of bits of the identification field one node of the subset of nodes when the subset of nodes includes fewer than a second predefined number of nodes.

17. The system of claim 16, wherein

10 the directory entry further includes a state field, the state field indicating whether the protocol engine is associating with a respective subset of bits of the identification field one node of the plurality of nodes or associating with each respective bit of the identification field one or more nodes of the plurality of nodes.

18. The system of claim 17, wherein

15 the protocol engine associates with a respective subset of bits of the identification field one node of the subset of nodes when the subset of nodes includes fewer than a second predefined number of nodes.

19. A protocol engine implementing a cache coherence protocol, for use in a multiprocessor computer system, the protocol engine located at a particular node of a plurality of nodes in the multiprocessor computer system, the protocol engine
20 comprising:

input logic for receiving a first invalidation request, the invalidation request identifying a memory line of information and including a pattern of bits for identifying a subset of the plurality of nodes that potentially store cached copies of the identified memory line; and

25 processing circuitry, responsive to receipt of the first invalidation request, for

sending a second invalidation request corresponding to the first invalidation request to a next node if the plurality of bits in fact identify the next node;

sending an invalidation acknowledgment to a requesting node identified in the first invalidation message if the plurality of bits fail to identify a next node; and

invalidating a cached copy of the identified memory line, if any, in the particular node of the plurality of nodes in the multiprocessor computer system.

20. A protocol engine implementing a cache coherence protocol, for use in a multiprocessor computer system, the protocol engine located at a particular node of a plurality of nodes in the multiprocessor computer system, the protocol engine comprising:

input logic for receiving a first invalidation request, the invalidation request identifying a memory line of information and including a pattern of bits for identifying a subset of the plurality of nodes that potentially store cached copies of the identified memory line; and

processing circuitry, responsive to receipt of the first invalidation request, for determining a next node identified by the pattern of bits in the invalidation request and for sending to the next node, if any, a second invalidation request corresponding to the first invalidation request, and for invalidating a cached copy of the identified memory line, if any, in the particular node of the multiprocessor computer system.

21. The protocol engine of claim 20, wherein the processing circuitry is configured to determine when the particular node is a last node identified by the pattern of bits in the invalidation request, and when said determination is made, to send an invalidation acknowledgment message to a requesting node identified in the first invalidation message.